



Practitioner's Docket No. 1956/135

PATENT

AF/2113
IFW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: James E. Dibb

Application No.: 09/879,554

Group No.: 2113

Filed: 06/12/2001

Examiner: S.T. Baderman

For: Method and System for Repairing a Redundant Array of Disk Drives

Mail Stop Appeal Briefs – Patents

Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-1450

TRANSMITTAL OF APPEAL BRIEF
(PATENT APPLICATION--37 C.F.R. § 41.37)

1. Transmitted herewith is the APPEAL BRIEF in this application, with respect to the Notice of Appeal filed on September 20, 2005.
2. STATUS OF APPLICANT

This application is on behalf of other than a small entity.

CERTIFICATION UNDER 37 C.F.R. §§ 1.8(a) and 1.10*

(When using *Express Mail*, the *Express Mail* label number is **mandatory**;
Express Mail certification is optional.)

I hereby certify that, on the date shown below, this correspondence is being:

MAILING

☒ deposited with the United States Postal Service in an envelope addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

37 C.F.R. § 1.8(a)

☒ with sufficient postage as first class mail.

37 C.F.R. § 1.10*

☐ as "Express Mail Post Office to Addressee"

Mailing Label No. _____ (mandatory)

TRANSMISSION

☐ facsimile transmitted to the Patent and Trademark Office, (703) _____ - _____

Signature

Date: November 9, 2005

John L. Conway

(type or print name of person certifying)

* Only the date of filing (§ 1.6) will be the date used in a patent term adjustment calculation, although the date on any certificate of mailing or transmission under § 1.8 continues to be taken into account in determining timeliness. See § 1.703(f). Consider "Express Mail Post Office to Addressee" (§ 1.10) or facsimile transmission (§ 1.6(d)) for the reply to be accorded the earliest possible filing date for patent term adjustment calculations.

3. FEE FOR FILING APPEAL BRIEF

Pursuant to 37 C.F.R. § 41.20(b)(2), the fee for filing the Appeal Brief is:

other than a small entity \$500.00

Appeal Brief fee due \$500.00

4. EXTENSION OF TERM

The proceedings herein are for a patent application and the provisions of 37 C.F.R. § 1.136 apply.

The total fee due is:

Appeal brief fee \$500.00

Extension fee (if any) \$0.00

TOTAL FEE DUE \$500.00

6. FEE PAYMENT

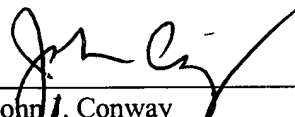
Attached is a check in the amount of \$500.00.

7. FEE DEFICIENCY

If any additional extension and/or fee is required, charge Deposit Account No. 19-4972.


A duplicate of this transmittal is attached.

Date: November 9, 2005



John L. Conway
Registration No. 48,241
Bromberg & Sunstein LLP
125 Summer Street
Boston, MA 02110-1618
U.S.
617-443-9292
Customer No. 02101

01956/00135 444166.1



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: James E. Dibb

Attorney Docket: 1956/135

Application No.: 09/879,554

Group Art Unit: 2113

Filed: June 12, 2001

Examiner: S. T. Baderman

For: Method and System for Repairing a Redundant Array of Disk Drives

CERTIFICATE OF MAILING

I hereby certify that this document, along with any other papers referred to as being attached or enclosed, is being deposited with the United States Postal Service as first class mail with sufficient postage in an envelope addressed to: Mail Stop Appeal Brief-Patents, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on November 9, 2005.

John L. Conway

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

Applicant submits this appeal brief for the above captioned application.

Table of Contents

Real Party in Interest..... 3

Related Appeals and Interferences..... 4

Status of Claims 5

Status of Amendments 6

Summary of Claimed Subject Matter 7

Grounds of Rejection to be Reviewed on Appeal..... 8

Argument 9

Claims Appendix 20

Evidence Appendix..... 23

Related Proceedings Appendix 24

Real Party in Interest

EMC Corporation

Related Appeals and Interferences

None

Status of Claims

Applicant appeals the final rejection of claims 1-14 and 16-19, which are pending in the application.

Status of Amendments

Cancellation of claim 15 subsequent to the final rejection is pending. The cancellation was filed on October 31, 2005.

Summary of Claimed Subject Matter

Note: all page and line references, except as noted, refer to the above-captioned U.S. patent application, no. 09/879,554 (“the application”).

The method of claim 1 restores the redundancy in a redundancy group of disk drives within a disk array when a drive within the group fails. A mirrored subsystem is created within the array. This subsystem includes the failed disk drive and a temporary disk drive (p. 6, lines 10-13). The redundancy group is reconfigured to consist of the working drives of the redundancy group plus the mirrored subsystem, thus restoring the redundancy of the redundancy group (p. 6, lines 14 to page 7, line 7). Even when the failed disk drive contains redundancy data for the redundancy group, the mirrored subsystem fully replaces the failed disk drive thereby fully restoring the group’s redundancy.

The computer program product of claim 16 contains program code that restores the redundancy in a redundancy group of disk drives within a disk array when a disk drive within the group fails. A mirrored subsystem is created within the array. The subsystem includes the failed disk and a temporary drive (p. 6, lines 10-13). The redundancy group is reconfigured to consist of the working drives of the redundancy group plus the mirrored subsystem such that the mirrored subsystem substitutes for failed disk drive in the redundancy group (p. 6, lines 14-20). The data blocks of the failed disk drive are reconstructed and written to the mirrored subsystem, restoring the redundancy of the redundancy group (p. 6, lines 21-27).

The disk drive array system of claim 18 includes logic to restore the redundancy of a redundancy group within the array when a disk drive of the group fails. A mirrored subsystem is created within the array. The subsystem includes the failed disk and a temporary drive (p. 6, lines 10-13). The redundancy group is reconfigured to consist of the working drives of the redundancy group plus the mirrored subsystem such that the mirrored subsystem substitutes for the failed disk drive in the redundancy group (p. 6, lines 14-19). The data blocks of the failed disk drive are reconstructed and written to the mirrored subsystem, restoring the redundancy of the redundancy group (p. 6, lines 21-27).

Grounds of Rejection to be Reviewed on Appeal

The rejections involved in this appeal are:

- 1) Claims 1-14 stand rejected under 35 U.S.C. § 112, first paragraph, as failing to comply with the written description requirement.
- 2) Claims 16-19 stand rejected under 35 U.S.C. § 102(e) as being anticipated by Kedem, U.S. pat. no. 6,154,853.

Argument

Issue 1: Rejection of Claims 1-14 under 35 U.S.C. § 112, first paragraph

Claims 1-14 stand rejected under 35 U.S.C. § 112, first paragraph, as failing to comply with the written description requirement. As will be shown below, the application contains an adequate written description of the claimed subject matter, as it would be understood by one skilled in the art. This rejection stems from a strained and illogical construction of the term “data block” in the specification.

The claimed invention repairs a failed disk drive in an array of disk drives. These arrays are often called a redundant array of inexpensive disks (“RAID”). The drives in the array are organized into redundancy groups. Redundancy groups allow the array to continue to store and supply data to users when one or more drives in the array fails. As the application states:


“For example, one common RAID configuration, called RAID-5, includes several disk drives, where data is written in a “stripe” across the drives, **consisting of corresponding data sectors on each drive, in which one of the data sectors stores parity data for the other data sectors in the stripe.** These several disk drives form a redundancy group. If one of the drives suffers a failure rendering its portion of a stripe inaccessible, the lost data is reconstructed from the other drives via computation methods well known in the art, such as employing an exclusive-or operation among the data sectors on the other drives” (emphasis added). (See, page 1, lines 13-21).

As the bolded words indicate, data sectors store data and the data includes redundancy data, in this case, parity data.

Another common redundancy group configuration is a mirrored subsystem or mirror. A “mirrored subsystem” is a disk drive array with slot connections for two disk drives. The subsystem includes logic such that data written to one disk drive is also copied to the other disk drive, when the other disk drive is connected to its slot. Data can be read from the subsystem by accessing either of the two disk drives if both drives are connected to their respective slots. The subsystem restores consistency of data as required between the two disk drives by copying data missing on one drive to the other drive. If

one drive fails, the configuration continues to operate without interruption, using the other disk drive. The subsystem restores consistency of data as required between the two disk drives by copying data missing on one drive to the other drive when the failed drive is replaced. (See, page 5, lines 9-16; see also, page 2, lines 4-7).)

The claimed embodiments of the present invention respond to failure of a disk drive within a redundancy group by creating a mirrored subsystem. For example, fig. 1 of the application shows a redundancy group 7 consisting of disk drives A, B, and C, (D is a spare drive):



When drive A fails, the drives are reconfigured as shown in fig. 4 of the application:




Fig. 4

Drives A and D form a mirrored subsystem, labeled “40.” The mirrored subsystem appears to the rest of the redundancy group as if it were the failed drive, except, of course, that the mirrored subsystem appears to be an operational drive (see page 3, lines 5-7). As the application states: “Each data block of the failed drive is then reconstructed (and)... written to the mirrored subsystem.” (See p. 3, lines 7-9.) Both user (logical) data and redundancy (e.g. parity) data are reconstructed to the mirrored subsystem. When reconstruction of the data on the failed drive to the mirrored subsystem is complete, the redundancy of the redundancy group is restored. This result follows inherently since the mirrored subsystem appears to the rest of the redundancy group as if it were the failed drive. (See page 3, lines 5-7). When the data on the failed drive has been reconstructed to the mirror, the redundancy group is again operational, except that the mirrored subsystem fills in for the failed disk drive. Advantageously, until physical repair of the failed drive can be accomplished, the redundancy group is restored to full function. For example, now either drive B or C could fail before the failed drive is physically replaced and the remaining drives could continue to supply and store data.

The examiner has proposed another construction for the teachings of the application. (See office action of 5/20/05, page 6, second paragraph). Rather than attaching the ordinary, customary meaning of “each data block” to mean “all blocks containing data”, the examiner has construed data blocks to mean only blocks containing logical data, i.e., the data generated by users of the disk array. As the office action states:

“The Applicant argues ‘a person of ordinary skill in the art would understand that the term ‘data block,’ as used in the specification, means all blocks on the failed disk drive, including both logical volumes and redundancy data.’ The Examiner respectfully disagrees. On page 6, lines 27-28, the specification specifically teaches that only ‘reconstructed data’ is stored on the mirrored subsystem. Not only does the Examiner disagree with the above statement that “a person of ordinary skill would understand.. .’ for there is no explanation as to ‘why’ one would understand this” - but the Examiner further believes that by teaching that the data stored on the mirrored subsystem is “reconstructed data,” **a person skilled in the art would have not understood this to include redundancy data - only reconstructed data.**” (Office action of 5/20/05, page 6, second paragraph, emphasis added).

As is apparent from the office action passage cited above, the examiner has overlooked the fact that when drive A fails, all data sectors on drive A must be reconstructed in order to be written to drive D. The reconstructed data sectors may contain logical (user) data or redundancy (parity) data. For both types of data sectors, the data is reconstructed from the data on drives B and C. As the specification of the present application states:

“Logic in array controller 25 reconstructs 130 each data block on failed drive A by accessing the data on remaining drives B and C. **The reconstruction is accomplished by error correction procedures that are well known in the art, such as performing an exclusive-or operation on the remaining disk drives that have not failed, if the error correction technique used is parity, or by using another technique appropriate for the error correction technique employed in the redundancy group in the array.** Logic in array controller 25 stores 140 each reconstructed data block to the mirrored subsystem 40.” (page 6, lines 21-28, emphasis added).

Clearly, if a data sector on drive A contains redundancy data, the data sector must still be reconstructed from data sectors on drive B and C.

According to the examiner's strained construction of the specification, only a portion of the data on the failed disk drive would be reconstructed to temporary drive D. The redundancy data for the redundancy group (e.g., parity) that resided on failed drive A would be ignored, even though temporary drive D has space for **all** of the data blocks on drive A. Under the examiner's strained construction, the redundancy group is not returned to full function until the failed drive is physically replaced, because the redundancy group does not contain enough redundancy data: the redundancy data on failed drive A has not been reconstructed. The illogical (and unintended) result is that formation of the mirrored subsystem would merely slow down drives B and C task of supplying all the data in the redundancy group. To this task would be added the task of reconstructing only the logical data on drive D, without restoration of the group's redundancy. Further, under this strained construction, if failed drive A contains only redundancy data, then no data would be written to drive D -- another unintended and illogical result.

The office action of 5/20/05, page 3, first paragraph, states: "Claim 1 includes the limitation: 'such that the mirrored subsystem is substituted for the failed disk drive in the redundancy group *and the redundancy of the redundancy group is restored, when the failed disk drive contains redundancy data for the redundancy group* (italicized words were amended into claim 1 on February 16, 2005).' After reviewing the specification thoroughly, the Examiner did not find any adequate written description for this amendment."


This rejection is overcome for claims 1-14 because a person of ordinary skill in the art would conclude that the inventor was in possession of the invention of claims 1-14 from the description in the specification. To satisfy the written description requirement, a patent specification must describe the claimed invention in sufficient detail that one skilled in the art can reasonably conclude that the inventor had possession of the claimed invention. See, e.g., Vas-Cath, Inc. v. Mahurkar, 935 F.2d at 1563, 19 USPQ2d at 1116; see also MPEP 2163.

A person skilled in the art of disk storage would reasonably conclude that the inventor of the present invention intended the ordinary meaning to attach to the words of the specification. The ordinary meaning of "each data block of the failed drive is then

reconstructed” is to reconstruct all data sectors on the failed drive, both data sectors containing logical data and sectors containing redundancy data. The patent application teaches using a mirrored subsystem to completely restore redundancy, even when the failed drive contained redundancy data. There is nothing in the application to suggest as the examiner asserts that redundancy is restored only when the failed drive contained only logical data. Instead, the application teaches when each data block, whether containing logical data or redundancy data, of the failed drive is reconstructed to the mirror, the redundancy group becomes whole again and, therefore, fully functional. To conclude otherwise would lead to the illogical result that the inventor intended the interim system (fig. 4 of the application, reproduced above) to be created for no advantage in data redundancy. Without a complete reconstruction, the interim configuration of fig. 4 would be open to catastrophic loss of data if drives B or C failed.

The examiner points to the Kedem reference (U.S. pat. no. 6,154,853) (“Kedem ‘853) as demonstrating that mirroring only the logical volumes of a failing drive is reasonable. (See office action of 5/20/05, page 7, first paragraph). In Kedem’s system, this approach makes sense only because the logical volumes on all of the drives of the redundancy group are copied to mirror drives rather than just the logical volumes on the single failing drive. (See Kedem ‘853, column 5, lines 9-20.)

Kedem’s disk drive array is organized as shown below in Kedem ‘853 fig. 3. Drives 32, 34, 36 and 38 form three redundancy groups (labeled 41, 43, and 45). D1 through D9 are logical volumes in the groups, while the shaded volumes contain redundancy data. (See Kedem ‘853, col. 4, lines 18 to 47).




Kedem '853 teaches restoring redundancy to a redundancy group, when a disk drive has failed, by mirroring all of the logical volumes of the entire redundancy group, not just the logical volumes of the single failing disk drive. If the redundancy group is formed on four disk drives, this approach requires three spare disk drives, (drives 31, 33, 35 in fig. 3). As Kedem puts it:

“.... device 34 has begun to fail and when that is sensed by the storage system, all the data from logical volume D2 of device 34 will be copied to a corresponding logical volume D2 of device 31. Additionally, all of the data from logical volume D5 of device 34 will be copied to a corresponding logical volume D5 of device 31. Since as described above, parity is not used in a mirroring scheme, the parity volume of device 34 will not be copied to spare device 31. Instead, the storage system will find the first logical volume of the next RAID group, here RAID group 45, and copy that data onto spare device 31. Thus for example, here data from logical volume D7 of device 32 will be copied to spare device 31. In order to provide complete protection for the data stored in the remaining devices 32, 36 and 38, **the remaining data from logical volumes D1, D4, D3, D6, D8 and D9 will be copied to the spare devices 33 and 35** (emphasis added).

After the mirror volumes have been created from the spare devices, the storage system will continue to operate in RAID mode while also maintaining mirror copies in the spare devices 31, 33 and 35. ... “(See Kedem '853, col. 5, lines 5-25).

This configuration is shown in Kedem fig. 6:



Kedem '853 restores redundancy with the help of $n-1$ temporary disk drives. Kedem's teachings are inapplicable to the description of the present invention in which a single spare drive restores redundancy. Given the architecture of the present invention, the specification teaches those of ordinary skill to reconstruct all data sectors (logical and redundancy) to the mirrored disk drive. Full redundancy is thus restored with just a single spare drive.

Since one skilled in the art would reasonably conclude that the inventor intended all data blocks on a failed drive to be reconstructed to the temporary drive in the mirrored subsystem, the inventor clearly had possession of the invention claimed in claims 1-14. Thus, the written description requirement of 35 U.S.C. §112, first paragraph, is satisfied for claims 1-14 and claims 1-14 should be allowed.

Issue 2: Rejection of Claims 16-19 under '102

Claims 16-19 stand rejected under 35 U.S.C. §102(e) as anticipated by Kedem, U.S. pat. no. 6,154,853.

Kedem '853 discloses a method of repairing a RAID storage subsystem. These subsystems include RAID groups with redundancy data (i.e., parity), such as a RAID-5 configuration. (See, for example, Kedem fig. 3, reproduced above, groups 41, 43, 45.) When a disk in a redundancy group begins to fail, the logical volumes on the failing disk in the RAID group are reconstructed to a spare disk drive. The redundancy data from the failing disk is not copied to the spare drive. (See, Kedem fig. 4 and col. 4, line 63 to col. 5, line 15.) The other logical volumes in the redundancy group are also copied to spare devices, forming a full mirror for each redundancy group's logical data. (See, Kedem '853, fig. 6 and col. 5, lines 17 to 21.) Once the failed drive has been replaced, this mirrored configuration, where only the logical volumes are mirrored, is converted back to a RAID group and the spare drives return to inactive status (see, Kedem '853 col. 5, lines 55-60). This conversion process includes recreating the redundancy data (i.e. parity) for the RAID group, when the failing disk drive includes redundancy data for the RAID group, since the mirror does not include such data. (See, generally, Kedem fig. 6 and col. 6.)

Claim 16 of the present application requires (in part):

“...reconstructing each data block of the failed disk drive in the redundancy group; and writing each reconstructed data block to the mirrored subsystem.”

During patent examination, the pending claims must be “given their broadest reasonable interpretation consistent with the specification.” In re Hyatt, 211 F.3d 1367, 1372, 54 USPQ2d 1664, 1667 (Fed. Cir. 2000). (See also MPEP 2111). The broadest reasonable interpretation of the claims must also be consistent with the interpretation that those skilled in the art would reach. In re Cortright, 165 F.3d 1353, 1359, 49 USPQ2d 1464, 1468 (Fed. Cir. 1999). The words of a claim must be given their “plain meaning” unless they are defined in the specification. (See MPEP 2111.01)

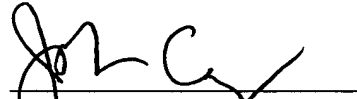
In the present application, , “data block” has been defined in the specification to mean all of the digital data on a disk drive and that is the only meaning that meets the objective of restoration of redundancy. Page 4, lines 25-30 of the specification defines the term “disk drive” as meaning, “.... a physical unit for storing randomly accessible **blocks of digital data**. A disk drive may be implemented, without limitation, as a rotating drive where data is stored electromagnetically on disk platters, or a solid state memory, such as flash memory, or any other physical means of **storing blocks of data that are randomly accessible** (emphasis added).” Both logical (user) data and redundancy data are included implicitly within the term “blocks of data” that are randomly accessible on the disk drive. In light of the definition of “disk drive” in the specification, “reconstructing each data block” must be construed in claim 16 as meaning reconstructing all of the data blocks, both logical and redundancy, of the failed disk drive. Further, as explained in issue 1 above, it would be inconsistent with the desire to restore redundancy to adopt the examiner’s strained construction of “each data block” in the claims. Kedem ‘853 teaches a different architecture, in which additional temporary drives are used such that only logical data need be copied. Kedem teaches away from reconstructing the redundancy data. Thus, Kedem ‘853 does not teach “...reconstructing each data block of the failed disk drive in the redundancy group; and writing each reconstructed data block to the mirrored subsystem” as required by claim 16 and Kedem ‘853 cannot anticipate claim 16. Claim 17 which depends from claim 16 and adds further limitations is, therefore, not anticipated by Kedem for at least the same reasons as for claim 16.

Likewise, claim 18 of the present application requires (in part):
“logic that reconstructs the data blocks on the failed drive to the mirrored subsystem.”
Since Kedem ‘853 does not teach reconstructing the redundancy data blocks of the failed disk drive to the mirrored subsystem, as shown above, Kedem cannot anticipate claim 18. Claim 19 which depends from claim 18 and adds further limitations is not anticipated by Kedem for at least the same reasons as for claim 18.

For all the foregoing reasons, Applicant submits that all claims in the application are allowable over the art of record and early notice to that effect is respectfully solicited.

DATE: November 9, 2005


Respectfully submitted,



John L. Conway.
Registration No. 48, 241
Attorney for Applicant

Bromberg & Sunstein LLP
125 Summer Street
Boston, MA 02110-1618
(617) 443-9292

01956/00135 426548.1



Claims Appendix

Application Serial No. 09/879,554

1. A method for handling a failed disk drive in a redundancy group of disk drives in an array of disk drives, the failed disk drive located in a failed disk drive slot, comprising:
creating a mirrored subsystem within the array, the subsystem including a temporary disk drive and the failed disk drive slot; and
reconfiguring the redundancy group to consist of the disk drives of the redundancy group that have not failed and the mirrored subsystem, such that the mirrored subsystem is substituted for the failed disk drive in the redundancy group and the redundancy of the redundancy group is restored, when the failed disk drive contains redundancy data for the redundancy group
2. A method as in claim 1, further comprising:
inserting a replacement disk drive in the failed disk drive slot;
copying data from the temporary disk drive to the replacement disk drive; and
replacing the mirrored subsystem with the replacement disk drive after the data on the replacement disk drive matches the data on the temporary disk drive.
3. A method as in claim 1, further comprising:
reconstructing each data block of the failed disk drive; and
writing each reconstructed data block to the mirrored subsystem.
4. A method as in claim 2, wherein the redundancy group is a RAID-5 array.
5. A method as in claim 2, wherein the redundancy group is a RAID-3 array.
6. A method as in claim 2, wherein the redundancy group is a RAID-1/0 array.
7. A method as in claim 2, wherein the redundancy group is a RAID-1 array.
8. A method as in claim 3, wherein the redundancy group is a RAID-5 array.
9. A method as in claim 3, wherein the redundancy group is a RAID-3 array.
10. A method as in claim 3, wherein the redundancy group is a RAID-1/0 array.
11. A method as in claim 3, wherein the redundancy group is a RAID-1 array.
12. A method as in claim 2, wherein the mirrored subsystem is a RAID-1 array.
13. A method as in claim 3, wherein the mirrored subsystem is a RAID-1 array.
14. A method as in claim 3, further comprising:

inserting a replacement disk drive in the failed disk drive slot;
copying data from the temporary disk drive to the replacement disk drive; and
replacing the mirrored subsystem with the replacement disk drive after the data on the replacement disk drive matches the data on the temporary disk drive.

15. cancellation pending.

16. A computer program product for use on a computer system for handling a failed disk drive in a redundancy group of disk drives in an array of disk drives, the failed disk drive located in a failed disk drive slot, the computer program product comprising a computer usable medium having computer readable program code thereon, the computer readable program code including program code for:

creating a mirrored subsystem within the array using a temporary disk drive and the failed disk drive slot;

reconfiguring the redundancy group to consist of the disk drives of the redundancy group that have not failed and the mirrored subsystem, such that the mirrored subsystem is substituted for the failed disk drive in the redundancy group; and

reconstructing each data block of the failed disk drive in the redundancy group;
and

writing each reconstructed data block to the mirrored subsystem.

17. A computer program product as in claim 16, further including program code for:

copying data from the temporary disk drive to a replacement disk drive in the failed disk drive slot; and

replacing the mirrored subsystem with the replacement disk drive after the data on the replacement disk drive matches the data on the temporary disk drive.

18. A disk drive array system comprising:

a redundancy group comprising at least two disk drives and associated disk drive slots;

a temporary disk drive with an associated temporary disk drive slot;

logic that detects a failure of one of the disk drives in the redundancy group;

logic that reconfigures the redundancy group to comprise the disk drives in the redundancy group that have not failed and a second storage array, the second storage array operating as a mirrored subsystem comprising the temporary disk drive and the disk drive slot associated with the failed disk; and
logic that reconstructs the data blocks on the failed drive to the mirrored subsystem.

19. A disk drive array system as in claim 18, further including:

logic that restores the redundancy group to its initial configuration, a replacement disk drive replacing the failed disk drive, after the temporary disk drive and the replacement drive inserted in the disk drive slot associated with the failed disk drive contain the same data.

Evidence Appendix

None

Related Proceedings Appendix

None.

426548.2